


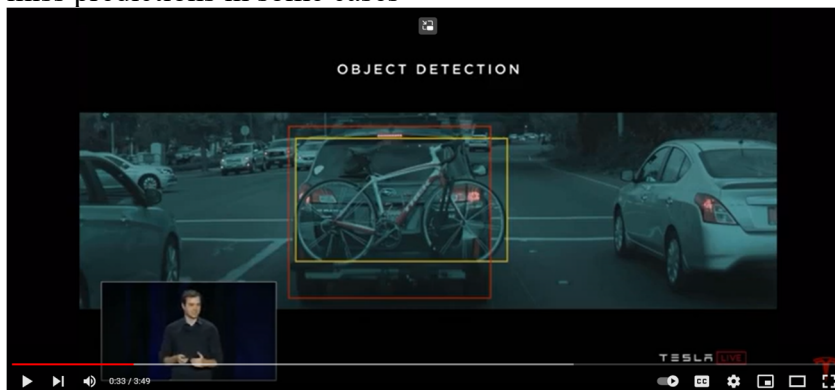
EXHIBIT J

U.S. Patent No. 11,238,344		
Claim 1	Exemplary Infringement Evidence ¹	
[1pre] A system comprising:	<p>Each autonomous Tesla vehicle with Software Version 9.0 and beyond (this includes vehicles with enhanced autopilot and/or full self-driving (FSD)) alone or together with the Dojo super computer infringes the '344 patent. Each of these vehicles is an example of a system that includes a processor, memory, etc. The Tesla fleet includes over a million of such customer vehicles. A fleet of Tesla vehicles alone or together with the Tesla DoJo super computer is also an example of a covered system.</p> <p>As a driver drives a Tesla vehicle, the vehicle learns various circumstances (e.g., pedestrians, other vehicles, roads, buildings, bikes, road debris, animals, etc.) as well as driving instructions (e.g., instructions for effecting speed, steering, breaking, trajectory, etc.) that the driver used to navigate those circumstances. This driving knowledge learned on one Tesla vehicle is transmitted to the Tesla DoJo computer and then distributed to all vehicles in Tesla fleet via over-the-air (OTA) software updates. The fleet, therefore, enables each vehicle in the fleet to autonomously implement driving instructions learned on one or more originating vehicles when similar circumstances are detected by any of the vehicles in Tesla fleet, thereby enabling their autonomous driving.</p> <p>Tesla explains its object detection (e.g., circumstance representation) technology as follows:</p> 	

¹ These infringement contentions are prepared with publicly available information.

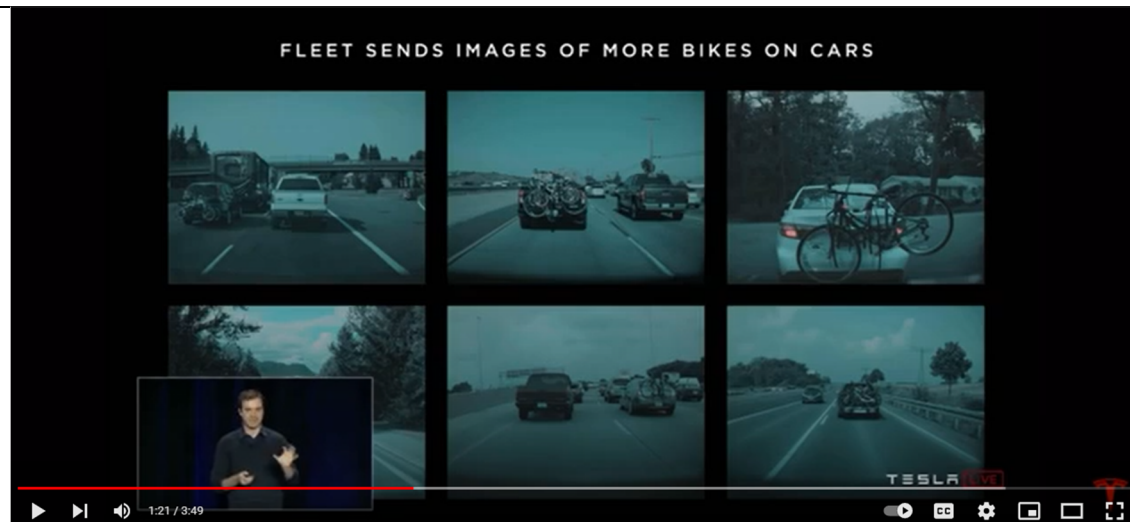
<https://www.youtube.com/watch?v=33K3id2xNAE&t>

“so object detection is something we care a lot about we’d like to put bounding boxes around say the cars and the objects here because we need to track them and we need to understand how they might move around so again we might ask human annotators to give us some annotations for these and humans might go in and might tell you that ok those patterns over there are cars and bicycles and so on and you can train your neural network on this but if you’re not careful the neural network all will make miss predictions in some cases”



<https://www.youtube.com/watch?v=33K3id2xNAE&t>

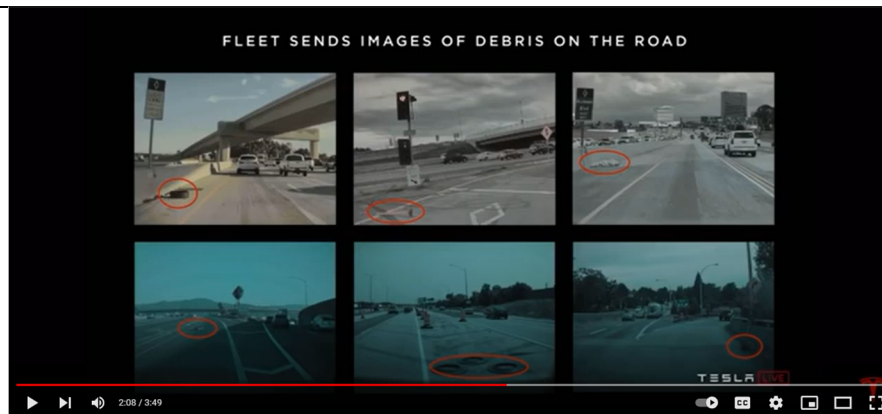
“so as an example if we stumble by a car like this that has a bike on the back of it then the neural network actually when I joined would actually create two deductions it would create a car deduction and a bicycle deduction and that’s actually kind of correct because I guess both of those objects actually exist but for the purposes of the controller and a planner downstream you really don’t want to deal with the fact that this bicycle can go with the car the truth is that that bike is attached to that car so in terms of like just objects on the road there’s a single object a single car and so what you’d like to do now is you’d like to just potentially annotate lots of those images as this is just a single car so the process that we that we go through internally in the team is that we take this image or a few images that show this pattern and *we have a mechanism a machine learning mechanism by which we can ask the fleet to source us examples that look like that and the fleet might respond with images that contains those patterns*”



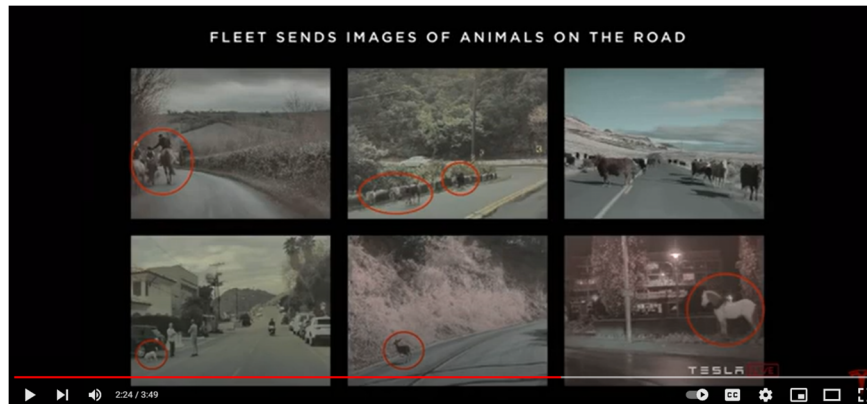
“so as an example these six images might come from the fleet they all contain bikes on backs of cars and we would go in and we would annotate all those as just a single car and then the the performance of that detector actually improves and the network internally understands that hey when the bike is just attached to the car that’s actually just a single car and it can learn that given enough examples and that’s how we sort of fix that problem ... now the fleet doesn’t just respond with bicycles on backs of cars we look for all the thing we look for lots of things all the time.”



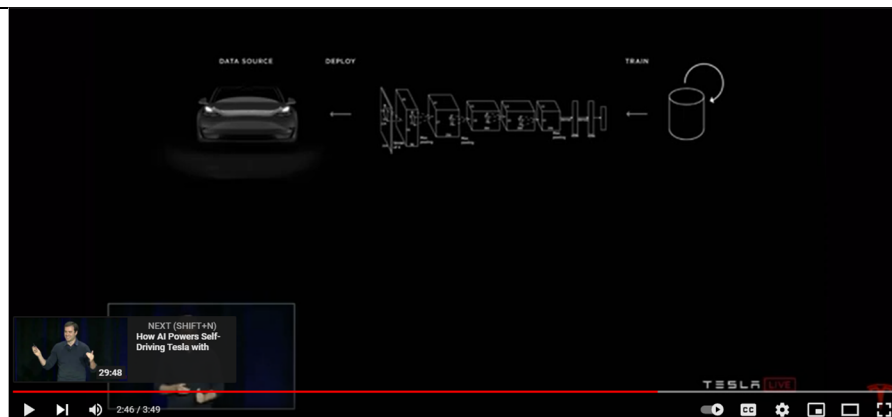
“so for example we look for boats and the fleet can respond with boats we look for construction sites and the fleet can send us lots of construction sites from across the world we look for even slightly more rare cases”



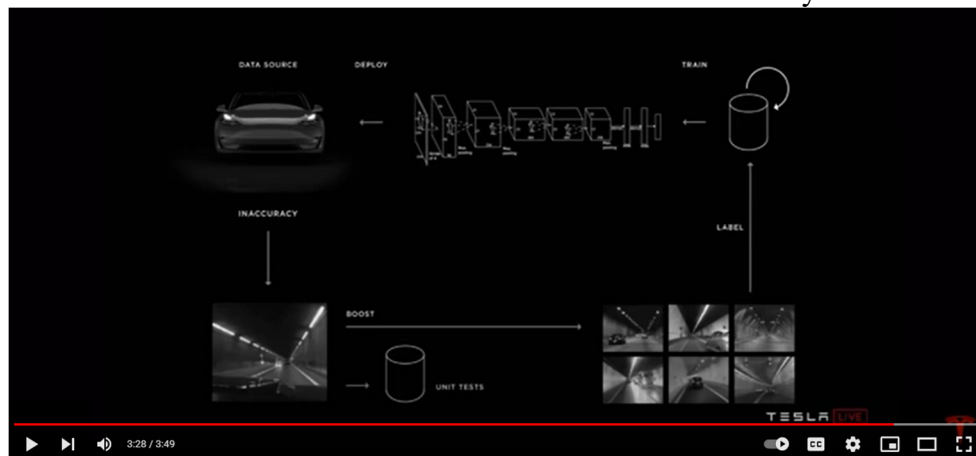
“so for example finding debris on the road is pretty important to us so these are examples of images that have streamed to us from the fleet that show tires cones, plastic bags and things like that if we can source these at scale we can annotate them correctly and *the neural network will learn how to deal with them in the world*”



“here’s another example animals of course also a very rare occurrence an event but we wanted the neural network to really understand what’s going on here that these are animals and we want to deal with that correctly”

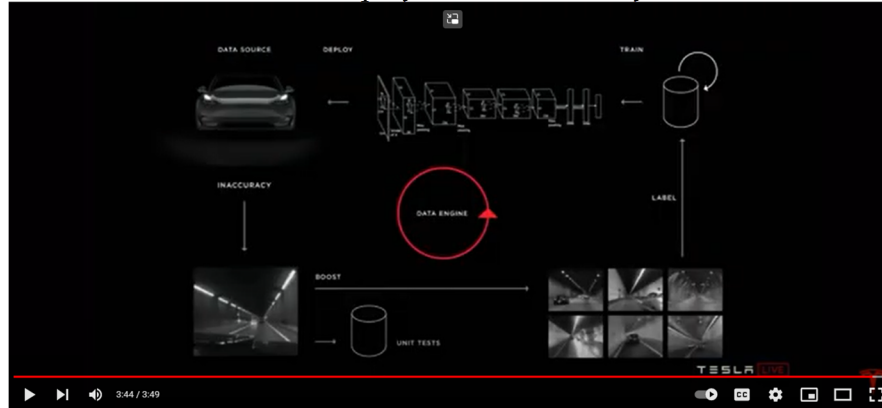


“so to summarize the process by which we iterate on neural network predictions looked something like this we start with a seed data set that was potentially sourced at random we annotate that data set and then we train your networks on that data set and put that in the car and then we have mechanisms by which we notice inaccuracies in the car when this detector may be misbehaving”



“so for example if we detect that the neural network might be uncertain or if we detect that or *if there’s a driver intervention or any of those settings we can create this trigger infrastructure that sends us data of those inaccuracies and so for example if we don’t perform very well on lane line detection on tunnels then we can notice that there’s a problem in tunnels that image would enter our unit tests so we can verify that we’ve actually fixing the problem over time but now what you do is to fix this inaccuracy you need to source many more examples that look like that so we asked the fleet to please send us many*

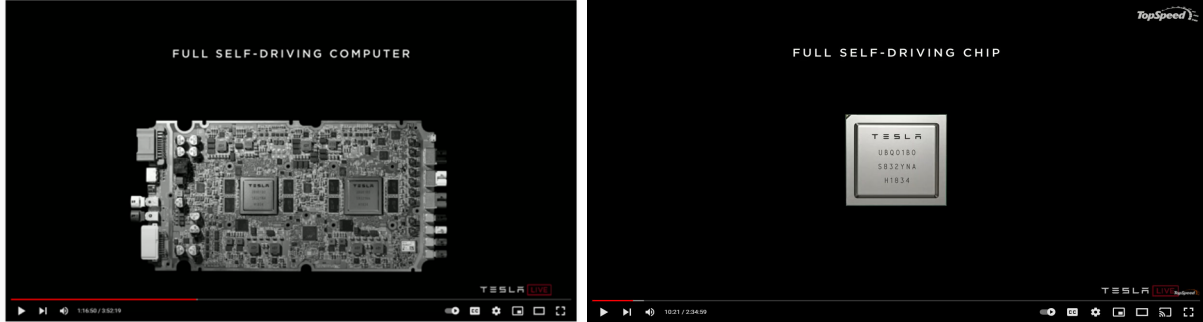
more tunnels and then we label all those tunnels correctly we incorporate that into the training set and we retrain the network redeploy and iterate the cycle over and over again and



“so we refer to this iterative process by which we improve these predictions as the data engine so iteratively deploying something potentially in shadow mode sourcing inaccuracies and incorporating the training set over and over again and we do this basically for all the predictions of these neural networks”

As Mr. Musk explained during Tesla AI Day all the human drivers are essentially training the neural net as to what is the correct course of action. See Tesla AI Day 2021 video <https://www.youtube.com/watch?v=j0z4FweCy4M> at 2:55:29. This is an example of the system being trained with object representations and instructions from a vehicle.

As Mr. Karpathy explained during Tesla Autonomy Day 2019 “while you are driving a car what you're actually doing is you are annotating the data because you are steering the wheel. you're telling us how to traverse different environments so what we're looking at here is a some person in the fleet who took a left through an intersection and what we do here is we have the full video of all the cameras and we know that the path that this person took because of the GPS, the inertial measurement unit, the wheel angle, the wheel ticks, so we put all that together and we understand the path that this person took through this environment and then of course we can use this for supervision [e.g., training a CNN through supervised learning] for the network so we just source a lot of this from the fleet, we train a neural network on those trajectories, and then the neural network predicts paths just from that data so really what ***this is referred to typically is called imitation learning we're taking human trajectories***

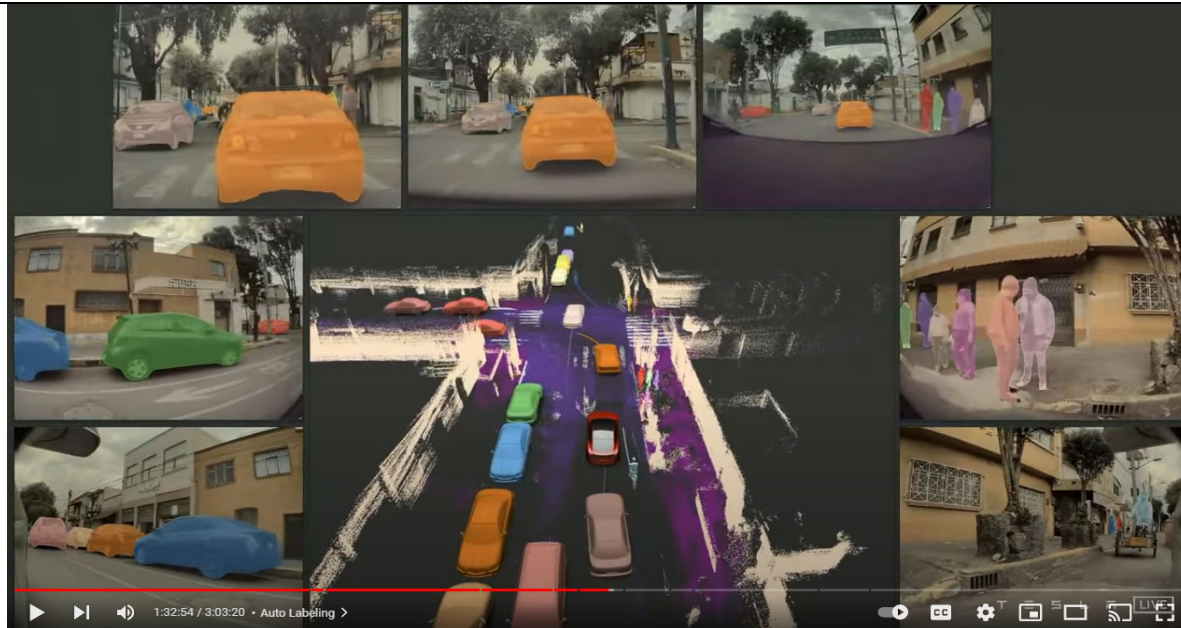
	<p><i>from the real world I'm just trying to imitate how people drive in real worlds and we can also apply the same data engine crank to all of this and make this work.”</i> https://www.youtube.com/watch?v=-b041NXGPZ8 at 1:04:10.</p> <p>In this way, driving knowledge is: (i) learned on a first Tesla vehicle, (ii) transferred to the DoJo super computer where it is synthesized with other driving knowledge, (iii) packaged and sent from the DoJo to the fleet, e.g., the driving knowledge is stored on a memory of a second Tesla vehicle via the fleet over-the-air (OTA) software update, and (iv) accessed on the second Tesla vehicle for its autonomous driving. The second Tesla vehicle can, thereby, implement the driving knowledge learned on the first Tesla vehicle and synthesized on the DoJo.</p>
<p>[1a] one or more processors configured to perform at least:</p>	<p>Each of the accused Tesla vehicles (Models 3, S, X, Y, etc.) includes one or more processors (e.g., the full self-driving computer/chip) configured to perform full self-driving.</p> <div data-bbox="571 638 1766 958">  </div> <p>See Tesla Autonomy Day 2019 video https://www.youtube.com/watch?v=-b041NXGPZ8 at 7:11 (Tesla full self driving computer) and at 10:22 (Tesla full self driving chip).</p>
<p>[1b] accessing a memory that stores at least a knowledgebase that includes a first circumstance representation correlated with a first one or more instruction sets for operating a first device,</p>	<p>The processor of the second Tesla vehicle (the claimed “second device”) accesses its memory that stores at least a knowledgebase that includes a first circumstance representation (e.g., representation of pedestrians, other vehicles, roads, buildings, etc.) correlated with a set of driving instructions (e.g., driving instructions for effecting speed, steering, breaking, trajectory, etc.; the claimed “first one or more instruction sets”) for operating a first Tesla vehicle (the claimed “first device”).</p>

wherein the first circumstance representation represents a first circumstance detected at least in part by one or more sensors of the first device, and

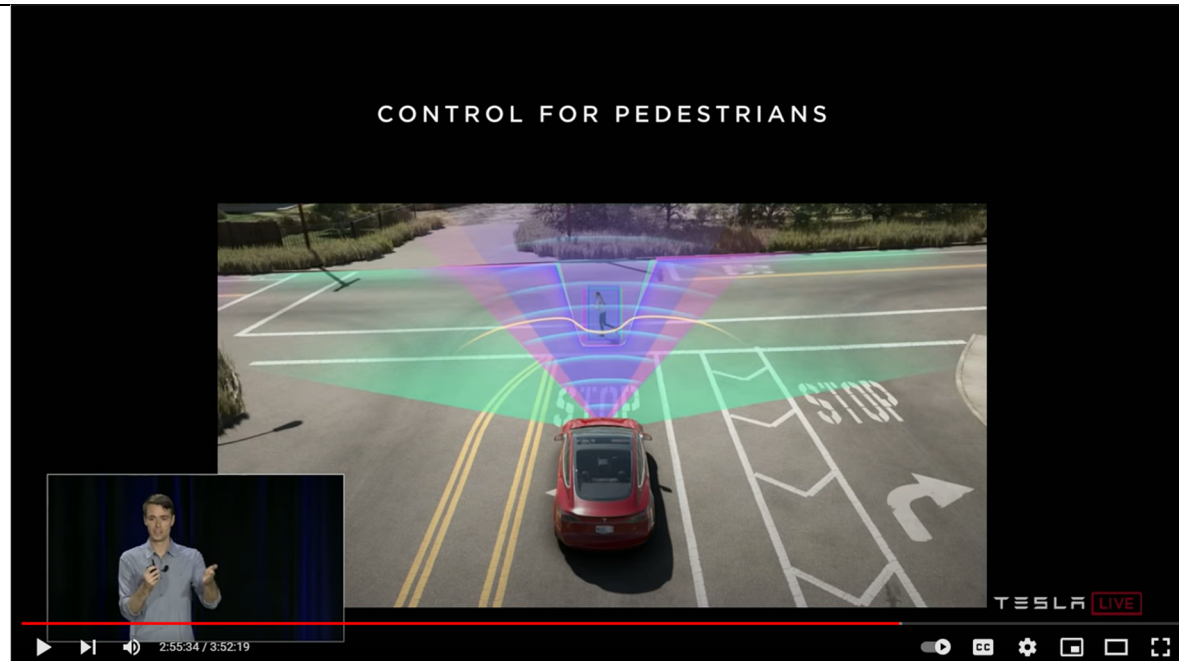
wherein at least a portion of the first circumstance representation or at least a portion of the first one or more instruction sets for operating the first device is learned in a learning process that includes operating the first device at least partially by a user;



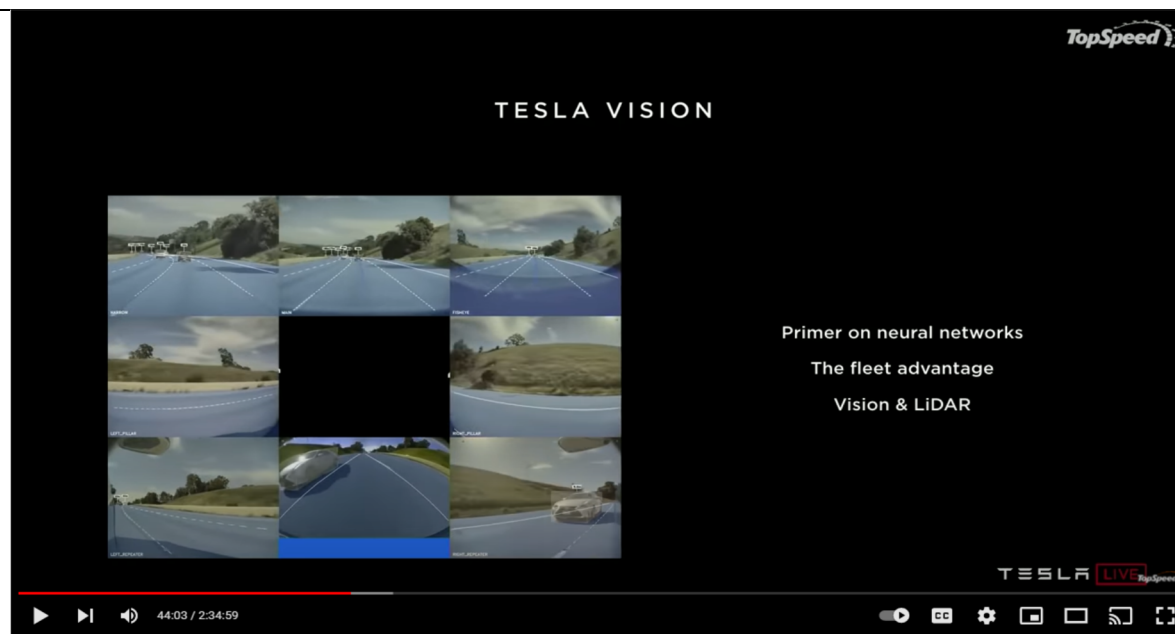
See Tesla AI Day 2021 video <https://www.youtube.com/watch?v=j0z4FweCy4M> at 1:31:58 (example of the claimed “circumstance representation”)



See Tesla AI Day 2021 video <https://www.youtube.com/watch?v=j0z4FweCy4M> at 1:32:40-48 (combining everything together we can produce these amazing data sets that annotate all of the road texture or the static objects and all of the moving objects [example of the claimed “circumstance representation”])

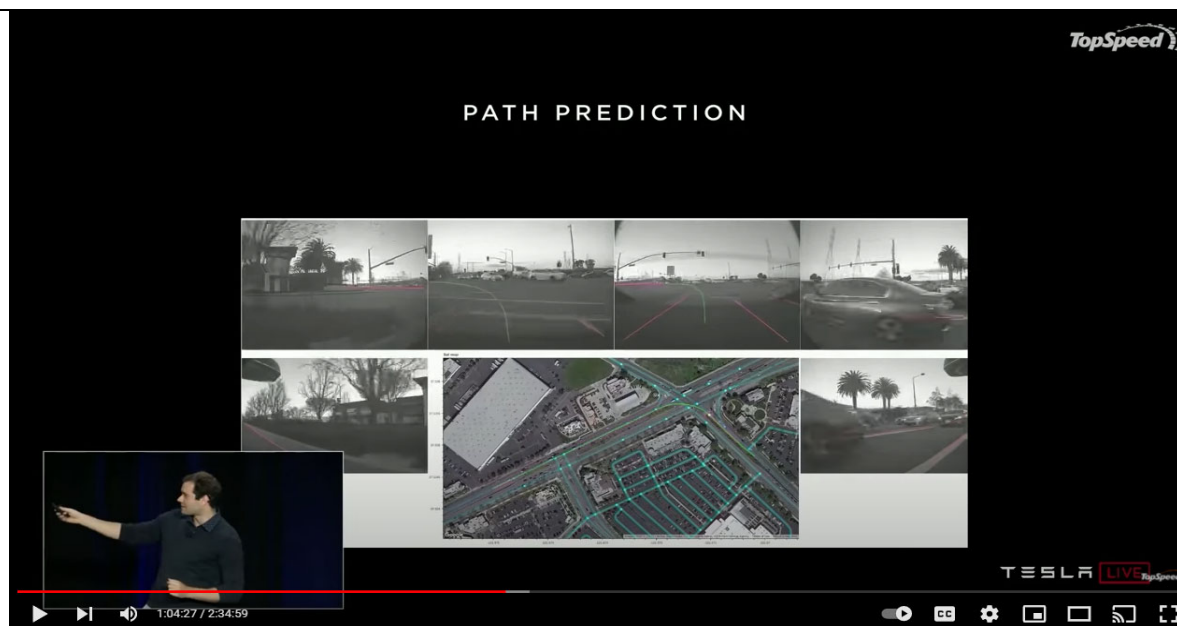


See Tesla AI Day 2021 video <https://www.youtube.com/watch?v=j0z4FweCy4M> at 2:55:34 (discussing an automatic emergency breaking system). As shown in the clip, each Tesla vehicle has sensors such as cameras, etc. that are used to detect a first circumstance, e.g., a pedestrian in front of the vehicle. The first circumstance representation, therefore, represents a circumstance such as a person in front of the vehicle (the claimed “wherein the first circumstance representation represents a first circumstance detected at least in part by one or more sensors of the first device”).



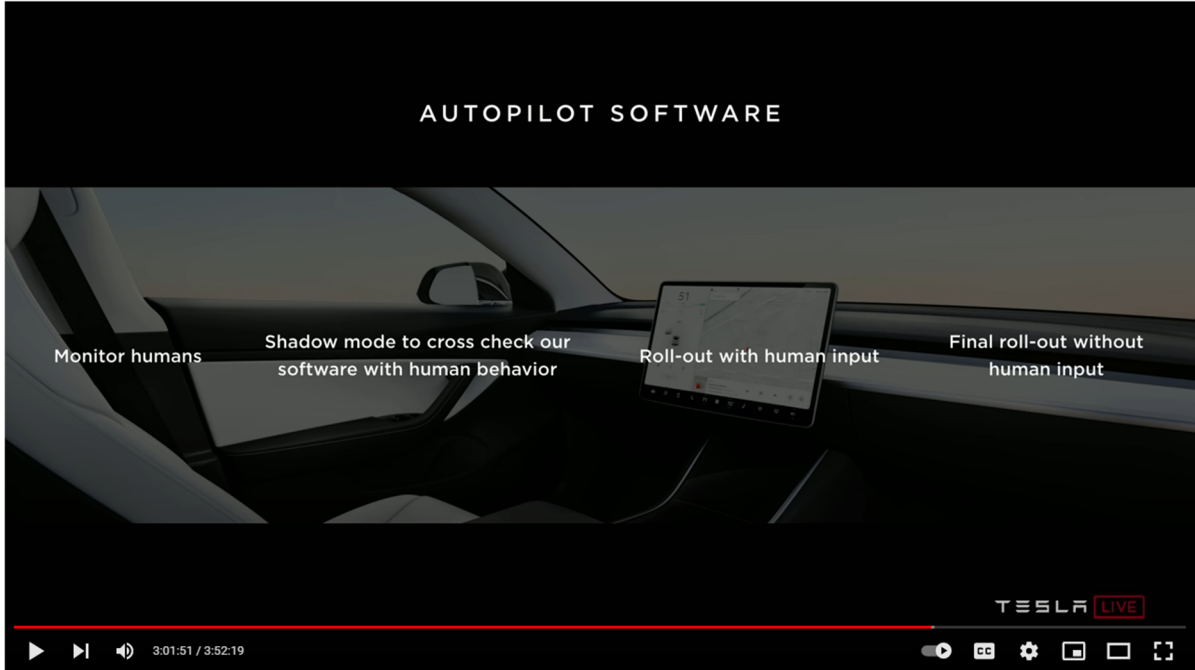
See also Tesla Autonomy Day 2019 video <https://www.youtube.com/watch?v=-b041NXGPZ8> at 44:00 (a stream of videos from eight cameras [the claimed “one or more sensors”] across the vehicle used to make a lane change).

As detailed above, the first circumstance representation (e.g., a representation of a pedestrian in front the vehicle, representation of surrounding vehicles in a lane change situation, animals or debris in the road, a stop light, a stop sign, etc.) and the corresponding driving instructions (e.g., instructions for applying the breaks so the pedestrian is not hit or so that traffic laws are obeyed, or instructions for turning the wheel to safely change lanes or avoid an animal or debris, etc.) are learned in a process that involves a driver (the claimed “user”) operating the first Tesla vehicle (the claimed “wherein at least a portion of the first circumstance representation or at least a portion of the first one or more instruction sets for operating the first device is learned in a learning process that includes operating the first device at least partially by a user”).

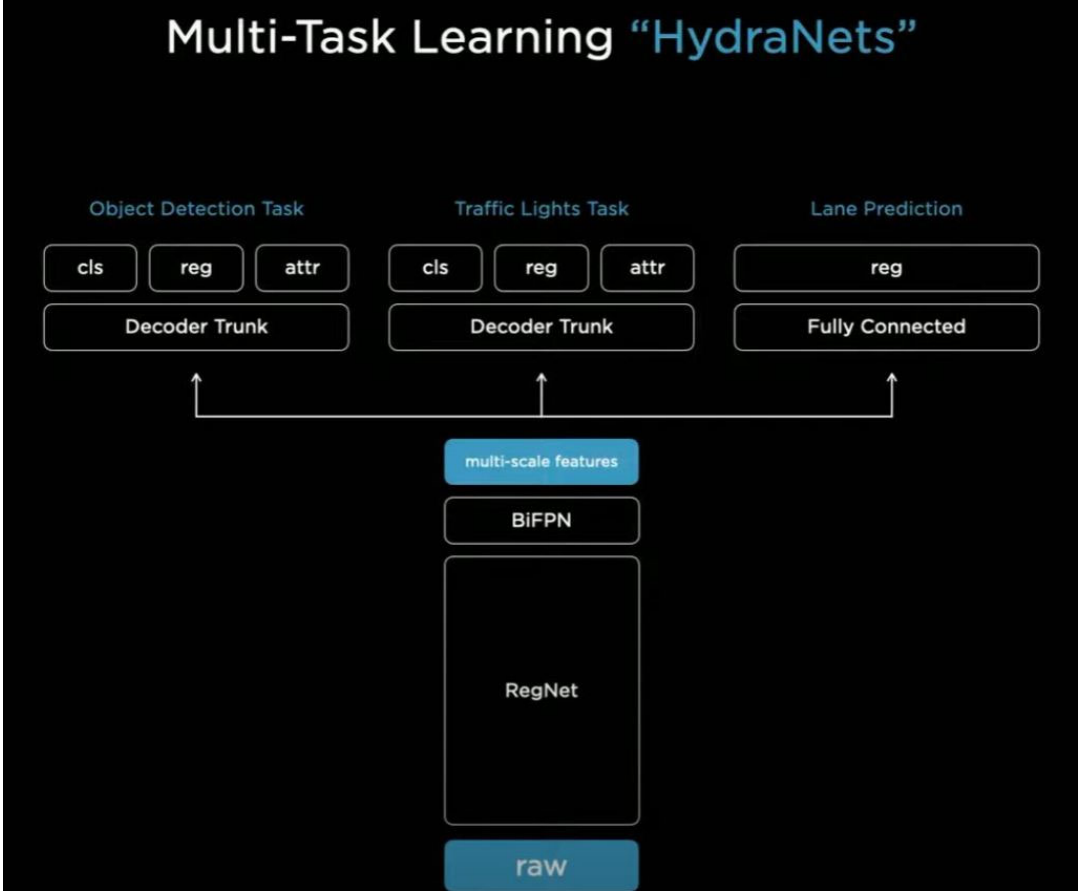


See Tesla Autonomy Day 2019 video <https://www.youtube.com/watch?v=-b041NXGPZ8> at 1:04:10 (while you [the claimed “user”] are driving a car what you’re actually doing is you are *annotating the data because you are steering the wheel* [e.g., **learned instruction**] you’re telling us how to traverse different environments so what we’re looking at here is some person in the fleet who took a left through an intersection and what we do here is we have the full video of all the cameras and we know that the path that this person took because of the GPS, the inertial measurement unit, the wheel angle, the wheel ticks, so we put all that together and we understand the path that this person took through this environment [**circumstance representation representing the environment correlated with driving instructions on how to traverse the environment**]. and then of course we can use this for supervision for the network so we just source a lot of this from the fleet, we train a neural network on those trajectories, and then the neural network predicts paths just from that data. ... we’re taking human trajectories from the real world we’re just trying to imitate how people drive in real worlds)

See Tesla AI Day 2021 video <https://www.youtube.com/watch?v=j0z4FweCy4M> at 2:55:29 (all the human drivers are essentially training the neural net as to what is the correct course of action [**the claimed driving instructions**])

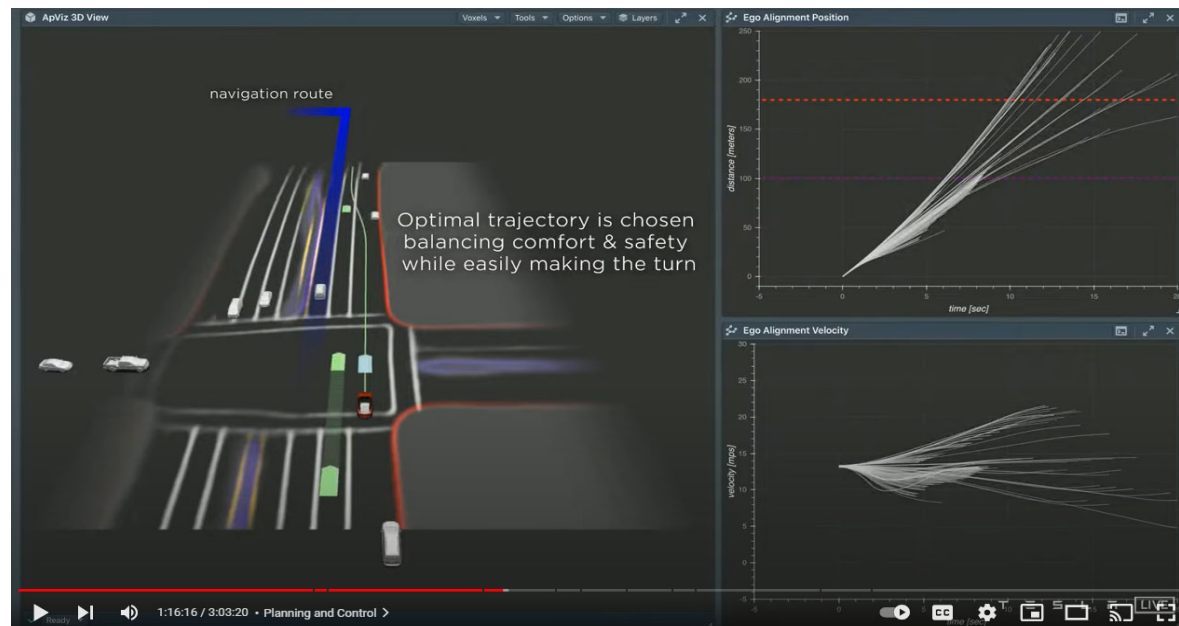
	 <p>See Musk, Karpathy, etc. in Tesla Autonomy Day 2019 video at https://www.youtube.com/watch?v=-b041NXGPZ8 at 1:03:50 (everyone is training the network all the time), 1:30:55 (Tesla system learns from drivers training the system), and 1:52:23 (high level chart clearly showing monitoring and learning from human drivers and implementing autonomous driving based on the learned knowledge).</p>
<p>[1c] generating or receiving a second circumstance representation, wherein the second circumstance representation represents a second circumstance detected at least in part by: the one or more sensors of the first device, <i>or one or</i></p>	<p>The processor of the second Tesla vehicle (the claimed “second device”) generates a second circumstance representation that represents a current circumstance (e.g., pedestrians, other vehicles, roads, buildings, animals, debris, etc.) detected at least in part by its cameras (the claimed “one or more sensors of a second device”).</p> <p>For instance, when the system is deployed to the fleet, a second Tesla vehicle (the claimed “second device”) generates a second circumstance representation (e.g., representation of a different pedestrian in front the vehicle, representation of different surrounding vehicles in a lane change situation, representation of a different intersection in a left turn situation, etc.) that is detected by the second Tesla vehicle’s sensors (e.g. cameras). The second circumstance representation is generated</p>

<p><i>more sensors of a second device;</i></p>	<p>so it can be compared with previously learned circumstance representations (the claimed “first circumstance representation”), and if a similar match (the claimed “at least partial match” in claim element 1d) is found, the correlated driving instructions (the claimed “first one or more instruction sets” in claim element 1b) can be used for autonomous driving.</p> <p>Analogous evidence as for “first circumstance representation” and “one or more sensors” recited in claim element 1b applies to “second circumstance representation” and “one or more sensors” here (e.g., information detected by the second car’s sensors), and it is not duplicated to save space. See claim element 1b.</p> <p>As explained during the 2021 AI day, Tesla said it “discovered that we don’t just want to detect cars, we want to do a large number of tasks. So for example, we want to do traffic light recognition and detection, a lane prediction and so on. So very quickly, we conversioned this kind of architectural layout, where there’s a common shared backbone, and then branches off into a number of heads. So we call these therefore Hydra Nets. And these are the heads of the Hydra.”</p> <p>https://elon-musk-interviews.com/2021/08/31/tesla-ai-day-the-presentation-i/</p>
--	--

	<h2 style="text-align: center;">Multi-Task Learning “HydraNets”</h2>  <p style="text-align: center;"> https://elon-musk-interviews.com/2021/08/31/tesla-ai-day-the-presentation-i/ </p>
<p>[1d] anticipating the first one or more instruction sets for operating the first device based on at least partial match between the second circumstance representation and the first circumstance</p>	<p>The processor of the second Tesla vehicle (the claimed “second device”) anticipates a set of driving instructions (e.g., instructions for applying the breaks so the pedestrian is not hit, instructions for turning the wheel to safely change lanes or to avoid an animal or debris, instructions for turning left in an intersection, etc.; the claimed “first one or more instruction sets”) based on similarity (the claimed “at least partial match”) between the first circumstance representation (e.g., a representation of a pedestrian in front the vehicle, a representation of surrounding vehicles in a lane change situation, a representation of an animal or debris, representation of an intersection in a left turn situation, etc. as previously learned on the first Tesla vehicle) and a second circumstance</p>

representation; and

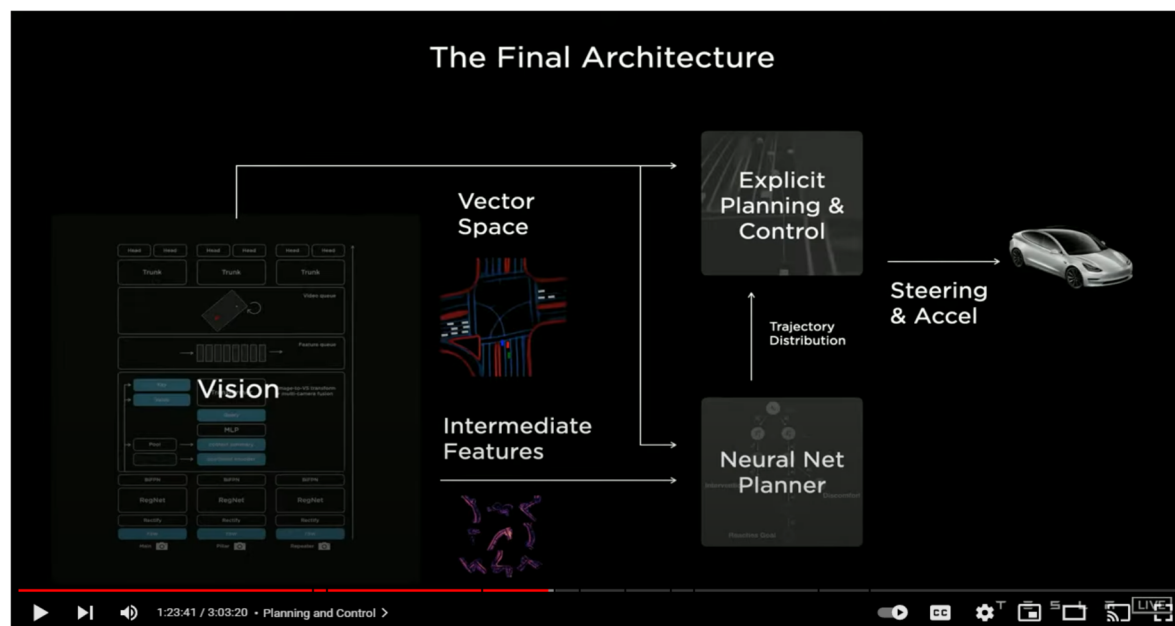
representation (e.g., a representation of a pedestrian in front the vehicle, a representation of surrounding vehicles in a lane change situation, a representation of an animal or debris, representation of an intersection in a left turn situation, etc. as currently detected by the second Tesla vehicle). Therefore, Tesla vehicles anticipate previously learned driving instructions based on similarity between the currently generated circumstance representation and previously learned circumstance representations.



See Tesla AI Day 2021 video <https://www.youtube.com/watch?v=j0z4FweCy4M> at 1:15:26 (let's see an example of how the search [e.g., search for previously learned driving instructions for effecting a trajectory correlated with a previously learned circumstance representation that at least partially matches the incoming circumstance representation] operates. so here we're trying to do a lane change. in this case the car needs to do two back to back lane changes to make the left turn ahead. for this, the car searches over different maneuvers. so the first one it searches is a lane change that's close by but the car breaks pretty harshly so it's pretty uncomfortable. the next maneuver it tries does the lane change a bit late so it speeds up, goes beyond the other car, goes in front of the other cars, and find it at the lane change but now it risks missing the left turn. we do thousands of such searches in

	<p>a very short time span). <i>See also id.</i> at 2:32:42 (similarly, for planning, we need to bake in a search [e.g., search for previously learned driving instructions for effecting a trajectory correlated with a previously learned circumstance representation that at least partially matches the incoming circumstance representation] and optimization into the planning, into the network architecture, and once we do that we should be able to do planning very quickly).</p>
<p>[1e] at least in response to the anticipating, executing the first one or more instruction sets for operating the first device, wherein the first device or the second device autonomously performs one or more operations defined by the first one or more instruction sets for operating the first device.</p>	<p>The processor of the second Tesla vehicle (the claimed “second device”) executes, at least in response to the anticipating (claim element 1d), the previously learned set of driving instructions (the claimed “first one or more instruction sets for operating the first device”) so that the second vehicle (the claimed “second device”) can drive autonomously based on the driving instructions learned on the first Tesla vehicle (the claimed “first device”).</p> <p>For instance, memory of the second Tesla vehicle stores a first circumstance representation (e.g., representation of a pedestrian in front the vehicle, representation of surrounding vehicles in a lane change situation, representation of an intersection in a left turn situation, etc.) correlated with a set of driving instructions (e.g., instructions for applying the breaks so the pedestrian is not hit, instructions for turning the wheel to safely change lanes, instructions for turning left in an intersection, etc.; the claimed “first one or more instruction sets”) that have previously been learned on a first Tesla vehicle (the claimed “first device”) and have been transferred to the second Tesla vehicle (the claimed “second device”) via the fleet over-the-air (OTA) software update. In response to determining that the incoming circumstance representation (e.g., representation of a pedestrian currently in front the vehicle, representation of currently surrounding vehicles in a lane change situation, representation of a current intersection in a left turn situation, etc.) is similar to (the claimed “at least partial match”) the previously learned circumstance representation (e.g., representation of a pedestrian previously in front the vehicle, representation of previously surrounding vehicles in a lane change situation, representation of a previous intersection in a left turn situation, etc.), the processor of the second Tesla vehicle (the claimed “second device”) causes the previously learned driving instruction (e.g., instructions for applying the breaks so the pedestrian is not hit, instructions for turning the wheel to safely change lanes, instructions for turning left in an intersection, etc.) to be executed.</p> <p>As stated by Elon Musk on Tesla Autonomy Day 2019, Tesla system distributes driving knowledge learned from multiple Tesla vehicles to all Tesla vehicles via over-the-air (OTA) software updates, thereby enabling the second Tesla vehicle (the claimed “second device”) to autonomously implement driving knowledge learned on the first vehicle (the claimed “first device”) as claimed. Mr. Musk and</p>

Mr. Karpathy describe this as a crucial competitive advantage of Tesla over other autonomous driving companies. See Musk, Karpathy, etc. in Tesla Autonomy Day 2019 video <https://www.youtube.com/watch?v=-b041NXGPZ8> at 38:58 (“I think a very powerful sustainable advantage for us is the fleet”), 44:40 (“it is such a big deal that we have the fleet”), 55:55 (“why is Tesla in such a unique and interesting position to really get all these three essentials right, and the answer to that of course is the fleet”), 122:14 (“it’s extremely difficult to catch up when Tesla has 100 times more miles per day than everyone else combined”).



See Tesla AI Day 2021 video <https://www.youtube.com/watch?v=j0z4FweCy4M> at 1:23:40 (so this is what a final architecture is going to look like. the vision system is going to crush down the dense video data into a vector space **[incoming circumstance representation representing objects in the second Tesla vehicle’s surrounding]**. it’s going to be consumed by both an explicit planner and a neural network planner **[CNN that stores previously learned circumstance representations correlated with previously learned driving instructions and that searches for at least partially matching circumstance representation and correlated driving instructions]** in addition to this, the network planner can also consume intermediate features of the network. Together, this produces a trajectory

	distribution and it can be optimized end to end both with explicit cost functions and human intervention and other imitation data. this then goes into explicit planning function that does whatever it sees for that and produces the final steering and acceleration commands for the car).
Claim 3	Exemplary Infringement Evidence
3. The system of claim 1, wherein the first circumstance representation includes a first one or more object representations, and wherein the second circumstance representation includes a second one or more object representations.	<p>Each autonomous Tesla vehicle with Software Version 9.0 and beyond (this includes vehicles with enhanced autopilot and/or full self-driving (FSD)) alone or together with the Dojo super computer includes the system of claim 1. Also, the first circumstance representation includes a first one or more object representations, and wherein the second circumstance representation includes a second one or more object representations.</p> <p>For example, the processor of the Tesla vehicle accesses its memory that stores at least a knowledgebase that includes a first circumstance representation including one or more object representations (e.g., representation of pedestrians, other vehicles, roads, buildings, etc.).</p> <p>And the second circumstance representation received or generated by the Tesla vehicle includes a second one or more object representations (e.g., representation of pedestrians, other vehicles, roads, buildings, etc.).</p>

